REMARKS

This Application has been carefully reviewed in light of the Office Action mailed August 3, 2004. Applicants appreciate the Examiner's consideration of the Application. Claims 2-5 and 23-25 have been canceled, and Claims 1, 6-9, 18, 22, 26-27, 31, 34, 38, 41, and 44-45 have been amended to clarify, more particularly point out, and more distinctly claim inventive concepts previously present in these claims. Applicants make no admission that these amendments narrow the scope of the claims or that the amendments are required for patentability. Applicants respectfully submit that no new matter has been added by the amendments to the claims. In order to advance prosecution of this Application, Applicants have responded to each notation by the Examiner. Applicants respectfully request reconsideration and favorable action in this case.

Objected to Claims

The Examiner states that Claims 18 and 34 are objected to as being dependent upon rejected base claims, but would be allowable if rewritten in independent form including all the limitations of the base claims and any intervening claims. (Office Action, page 11, paragraph 1.) Claims 18 and 34 have been amended to include the limitations of the base claims and any intervening claims. Accordingly, Applicants respectfully request allowance of Claims 18 and 34.

Section 102 Rejection

The Examiner rejects Claims 1-11, 13, 19-27, 29, 30, 35-38, 40, 41, and 43-45 under 35 U.S.C. § 102(a) as being unpatentable over "World Wide Web Caching: Trends and Techniques" by Barish et al. ("*Barish*"). Applicants respectfully traverse this rejection for the reasons discussed below.

Applicants respectfully submit that *Barish* fails to disclose, or even teach or suggest, the elements specifically recited in Applicants' claims. For example, *Barish* fails to disclose, or even teach or suggest, the following elements (1) and (2):

(1) establishing a primary distribution of the plurality of cache shares using the locator identifiers, the primary distribution indicating a first allocation of the plurality of cache shares among a plurality of clients; and

(2) establishing a secondary distribution of the plurality of cache shares using the locator identifiers, the secondary distribution indicating a second allocation of the plurality of cache shares among the plurality of clients to be used in place of the primary distribution in response to a trigger occurrence (recited in Claim 1, as amended).

Barish discloses a hierarchical caching technique. According to *Barish*:

> The basic idea is to have a series of caches hierarchically arranged in a tree-like structure and to allow those caches to leverage from each other when an object request arrives and the receiving cache does not have that object.
>
> Usually, in hierarchical designs child caches can query parent caches and children can query each other, but parents never query children. This promotes an architecture where information gradually filters down to the leaves of the hierarchy. In a sense, the adaptive caching approach also uses cache hierarchies (in the form of cache groups) to diffuse information from dynamic hot spots to the outlying cache clusters, but these hierarchies are peer-based: the parent/child relationships are established per information object. Thus, in one case a cache group might act as a parent for a set of information object X, but also as a child (or intermediary) node for information object Y.

(*Barish*, page 181, column 1, paragraph 6-column 2, paragraph 1.)

That is, *Barish* discloses a hierarchical caching technique for finding cached content in response to receiving a request for the content. The hierarchical caching technique, however, does not apply to allocating cached content. Accordingly, the hierarchical caching technique fails to disclose, teach, or suggest "establishing a primary distribution of the plurality of cache shares using the locator identifiers, the primary distribution indicating a first allocation of the plurality of cache shares among a plurality of clients," and "establishing a secondary distribution of the plurality of cache shares using the locator identifiers, the secondary distribution indicating a second allocation of the plurality of cache shares among the plurality of clients to be used in place of the primary distribution in response to a trigger occurrence," recited in amended Claim 1.

Barish also discloses a Cache Group Management Protocol. According to *Barish*:

> Adaptive caching uses the Cache Group Management Protocol (CGMP) and Content Routing Protocol (CRP). CGMP specifies how meshes are formed, and how individual caches join and leave those meshes. In general, caches are organized into overlapping multicast groups which use voting and feedback techniques to estimate the usefulness of admitting or excluding members from that group. The ongoing negotiation of mesh formation and membership results in a virtual topology.

The Cache Group Management Protocol of *Barish*, however, fails to disclose, teach, or suggest "establishing a primary distribution of the plurality of cache shares using the locator identifiers, the primary distribution indicating a first allocation of the plurality of cache shares among a plurality of clients," and "establishing a secondary distribution of the plurality of cache shares using the locator identifiers, the secondary distribution indicating a second allocation of the plurality of cache shares among the plurality of clients to be used in place of the primary distribution in response to a trigger occurrence," recited in amended Claim 1.

Consequently, at a minimum, *Barish* fails to disclose, or even teach or suggest, the elements specifically recited in independent Claim 1, as amended. Applicants respectfully request reconsideration and allowance of independent Claim 1. Independent Claims 22, 38, 41, 44, and 45 recite certain limitations substantially similar to those recited in independent Claim 1. Accordingly, for at least the same reasons, Applicants also respectfully request reconsideration and allowance of independent Claims 22, 38, 41, 44, and 45.

Applicants' dependent claims are allowable based on their dependence on the independent claim and further because they recite numerous additional patentable distinctions over the reference of the rejection. Because Applicants believe they have amply demonstrated the allowability of the independent claim over the reference of the rejection, and to avoid burdening the record, Applicants have not provided detailed remarks concerning these dependent claims. Applicants, however, remain ready to provide such remarks if it becomes appropriate to do so. Applicants respectfully request reconsideration and allowance all claims that depend on Claims 1, 22, 38, and 41.

Section 103(a) Rejection

The Examiner rejects under 35 U.S.C. § 103(a): Claims 12, 14, 28, 39, and 42 as being unpatentable over *Barish* in light of U.S. Patent No. 6,112,279 to Wang ("*Wang*"); and Claims 15-17 and 31-33 as being unpatentable over *Barish* in light of U.S. Patent No. 6,542,967 to Major ("*Major*"). Applicants respectfully traverse this rejection for the reasons discussed below.

Applicants respectfully submit that the combination of *Barish* and *Wang* as suggested by the Examiner fails to disclose, teach, or suggest elements specifically recited in Applicants' claims. For example, the *Barish-Wang* combination suggested by the Examiner

fails to disclose, teach, or suggest elements (1) and (2) as recited in independent Claims 1, 22, 38, and 41, as amended. As discussed above, *Barish* fails to disclose, teach, or suggest elements (1) and (2).

Moreover, *Wang* also fails to disclose, teach, or suggest the elements. *Wang* discloses:

> In an arrangement where a plurality of cache servers are interconnected to form a virtual cache, each cache server includes a selection module and a caching module. The selection module determines whether the cache server can service an incoming request for information, whether the request ought to be directed to another one of the cache servers, or whether the request be routed to the site from whence the information is requested.

(*Wang*, Abstract.) That is, *Wang* discloses a technique for determining whether a cache server can service a request. *Wang*, however, fails to disclose, teach, or suggest "establishing a primary distribution of the plurality of cache shares using the locator identifiers, the primary distribution indicating a first allocation of the plurality of cache shares among a plurality of clients," and "establishing a secondary distribution of the plurality of cache shares using the locator identifiers, the secondary distribution indicating a second allocation of the plurality of cache shares among the plurality of clients to be used in place of the primary distribution in response to a trigger occurrence." Thus, the *Barish-Wang* combination suggested by the Examiner fails to disclose, teach, or suggest elements (1) and (2) as recited in independent Claims 1, 22, 38, and 41, as amended.

Applicants' dependent claims are allowable based on their dependence on the independent claims and further because they recite numerous additional patentable distinctions over the references of the rejection. Because Applicants believe they have amply demonstrated the allowability of the independent claims over the references of the rejection, and to avoid burdening the record, Applicants have not provided detailed remarks concerning these dependent claims. Applicants, however, remain ready to provide such remarks if it becomes appropriate to do so. Accordingly, for at least the same reasons, Applicants also respectfully request reconsideration and allowance of Claims 12, 14, 28, 39, and 42.

Applicants respectfully submit that the combination of *Barish* and *Major* as suggested by the Examiner also fails to disclose, teach, or suggest elements specifically recited in Applicants' claims. For example, the *Barish-Major* combination suggested by the Examiner fails to disclose, teach, or suggest elements (1) and (2) as recited in independent Claims 1 and

22, as amended. As discussed above, *Barish* fails to disclose, teach, or suggest elements (1) and (2).

Moreover, *Major* fails to disclose, teach, or suggest the elements. According to *Major*:

> The invention comprises a cache object store organized to provide fast and efficient storage of data as cache objects, which can be organized into cache object groups. The cache object store preferably embodies a multi-level hierarchical storage architecture comprising (i) a primary memory-level (RAM) cache store and (ii) a secondary disk-level cache store, each of which is configured to optimize access to the cache object-groups. These levels of the cache object store further cooperate to provide an enhanced caching system that exploits persistent and non-persistent storage characteristics of the inventive architecture.

(*Major*, column 2, line 63-column 3, line 6.) *Major* also discloses:

> A cache object manager implements various aging and storage management algorithms to manage the cache object store. An example of such an aging policy is a modified least recently used (LRU) algorithm that strives to keep those object groups that are accessed most often in the primary-level cache store, with as many remaining object groups stored on the secondary-level store for quick retrieval.

(*Major*, column 2, lines 19-25.)

That is, *Major* discloses storing object groups that are accessed most often in a primary-level cache store and storing remaining object groups in a secondary-level cache store according to aging and storage management algorithms. The cache objects, however, are not stored according to locator identifiers. Moreover, *Major* discloses performing storage of object groups on a dynamic basis, so presumably the distribution of the object groups is not established beforehand. Accordingly, *Major*, fails to disclose, teach, or suggest elements "establishing a primary distribution of the plurality of cache shares using the locator identifiers, the primary distribution indicating a first allocation of the plurality of cache shares among a plurality of clients," and "establishing a secondary distribution of the plurality of cache shares using the locator identifiers, the secondary distribution indicating a second allocation of the plurality of cache shares among the plurality of clients to be used in place of the primary distribution in response to a trigger occurrence." Thus, the *Barish-Major* combination suggested by the Examiner fails to disclose, teach, or suggest elements (1) and (2) as recited in independent Claims 1 and 22, as amended.

Applicants' dependent claims are allowable based on their dependence on the independent claims and further because they recite numerous additional patentable distinctions over the references of the rejection. Because Applicants believe they have amply demonstrated the allowability of the independent claims over the references of the rejection, and to avoid burdening the record, Applicants have not provided detailed remarks concerning these dependent claims. Applicants, however, remain ready to provide such remarks if it becomes appropriate to do so. Accordingly, for at least the same reasons, Applicants also respectfully request reconsideration and allowance of Claims 15-17 and 31-33.

CONCLUSION

Applicants have made an earnest attempt to place this case in condition for allowance. For at least the foregoing reasons, Applicants respectfully request full allowance of all the pending claims.

If the Examiner believes a telephone conference would advance prosecution of this case in any way, the Examiner is invited to contact Keiko Ichiye, the Attorney for Applicants, at the Examiner's convenience at (214) 953-6494.

A check in the amount of $88.00 is enclosed as fees for two additional independent claims. Although Applicants believe no additional fees are due, the Commissioner is hereby authorized to charge any fees or credit any overpayments to Deposit Account No. 02-0384 of Baker Botts L.L.P.

Respectfully submitted,

BAKER BOTTS L.L.P.
Attorneys for Applicants

Keiko Ichiye
Reg. No. 45,460

KI/ls

Correspondence Address:
Baker Botts L.L.P.
2001 Ross Avenue, Suite 600
Dallas, Texas 75201-2980
(214) 953-6494

Date: 08/26/2004

Customer Number: 05073